

Computing Equilibria in Anonymous Games

Constantinos Daskalakis

Christos Papadimitriou *

University of California, Berkeley

Computer Science

{costis, christos}@cs.berkeley.edu

Abstract

We present efficient approximation algorithms for finding Nash equilibria in anonymous games, that is, games in which the players utilities, though different, do not differentiate between other players. Our results pertain to such games with many players but few strategies. We show that any such game has an approximate pure Nash equilibrium, computable in polynomial time, with approximation $O(s^2\lambda)$, where s is the number of strategies and λ is the Lipschitz constant of the utilities. Finally, we show that there is a PTAS for finding an ϵ -approximate Nash equilibrium when the number of strategies is two.

1 Introduction

Will you come to FOCS? This decision depends on many factors, but one of them is *how many other theoreticians will come*. Now, whether each of them will come depends in a small way on what you will do, and hence this aspect of the decision to go to FOCS is game-theoretic — and in fact of a particular specialized sort explored in this paper: Each player has a small number of strategies (in this example, two), and the utility of each player depends on her/his own decision, as well as on how many other players will choose each of these strategies. It is crucial that the utilities do *not* depend on the identity of the players making these choices (that is, we ignore here your interest in whether your friends will come). Such games are called *anonymous games*, and in this paper we give two polynomial algorithms for computing approximate equilibria in these games. In fact, our algorithms work in a generalized framework: The players can be divided into a few *types* (e.g., colleagues, students, big shots, etc.), and your utility depends on how many of the players of *each type* choose each of the

strategies.

Notice that this is a much more general framework than that of symmetric games (where all players are identical); each player can have a very individual way of evaluating the situation, and her/his utility can depend on the choices of others in an individual arbitrary way; in particular, there may be no monotonicity: For example, a player may prefer a mob with 1000 attendees, mostly students, to a tiny workshop of 20, while a medium-sized conference of 200 may be more attractive than either; a second player may order these in the exact opposite way. Anonymous games comprise a broad and well studied class of games (see e.g. [4, 5, 16] for recent work on this subject by economists) which are of special interest to the Algorithmic Game Theory community, as they capture important aspects of auctions and markets, as well as of Internet congestion.

Our interest lies in computing *Nash equilibria* in such games. The problem of computing Nash equilibria in a game was recently shown to be PPAD-complete in general [9], even for the two-player case [6]. Since that negative result, the research effort in this area was, quite predictably, directed towards two goals: (1) computing approximate equilibria (mixed strategy profiles from which no player has incentive *more than* ϵ to defect), and (2) exploring the algorithmic properties of special cases of games. The approximation front has been especially fertile, with several positive and negative results shown recently [19, 7, 18, 10, 11, 13].

What is known about special cases of the Nash equilibrium problem? Several important cases are now known to be generic; these include, beyond the aforementioned 2-player games, win-lose games (games with 0-1 utilities) [1], and several kinds of succinctly representable games such as graphical games [9] and anonymous games (actually, the even more specialized *symmetric* games [14]). For anonymous games, the genericity argument goes as follows: Any game can be made anonymous by expanding the strategy space so that each player first chooses an identity (and is punished if s/he fails to choose her/his own) and then a strategy; it is easy

*The authors were supported through NSF grant CCF - 0635319, a gift from Yahoo! Research and a MICRO grant.

to see that, in this expanded strategy space, the utilities can be rendered in the anonymous fashion. Note, however, that this manoeuvre requires a large strategy space; in contrast, for other succinct games such as the graphical ones, genericity persists even when the number of strategies is two [15]. Are anonymous games easier when the number of strategies is fixed? We shall see that this is indeed the case.

How about tractable special cases? Here there is a relative poverty of results. The zero-sum two-player case is, of course, well known [22]. It was generalized in [17] to low-rank games (the matrix $A + B$ is not quite zero, but has fixed rank), a case in which a PTAS for the Nash equilibrium problem is possible. It was also known that symmetric games with (about logarithmically) few strategies per player can be solved exactly in polynomial time by a reduction to the theory of real closed fields [23]. For congestion games we can find in polynomial time a pure Nash equilibrium if the game is a symmetric network congestion game [12], and an approximate pure Nash equilibrium if the congestion game is symmetric (but not necessarily network) and the utilities are somehow “continuous” [8]. Finally, in [20] Milchtaich showed that anonymous congestion games in graphs consisting of parallel links (equivalently, anonymous games in which the utility of a player, for each choice made by the player, is a nondecreasing function of the number of players who have chosen the same strategy) have pure Nash equilibria which can be computed in polynomial time by a natural greedy algorithm.

In this paper we prove two positive approximation results for anonymous games with a fixed number s of strategies. Our first result states that any such game has a pure Nash equilibrium that is ϵ -approximate, where ϵ is bounded from above by a function of the form $f(s)\lambda$. Here λ is the *Lipschitz constant* of the utility functions (a measure of continuity of the utility functions of the players, assumed to be such that for any partitions x and y of the players into the s strategies, $|u(x) - u(y)| \leq \lambda \|x - y\|_1$). To get a sense of scale for λ note that the arguments of u range from 0 to n and so, if u were a linear function in the range $[0, 1]$, λ would be at most $\frac{1}{n}$. $f(s)$ is a quadratic function of the number of strategies. That ϵ cannot be smaller than λ is easy to see (the matching pennies problem provides an easy example); the results of [8] for congestion games show a similar dependence on λ (what they call “the bounded jump property”). We conjecture that the dependence on s can be improved to $s\lambda$. Our proof uses Brouwer’s fixed point theorem on an interpolation of the (discrete) best-response function to identify a simplex of pure strategy profiles and from that produce, by a geometric argument, a pure strategy profile that is ϵ -approximate, with ϵ bounded as above.

Our second result is a PTAS for the case of two strategies. The main idea is to round the mixed strategies of the players to some nearby multiple of ϵ ; then each such quantized mixed strategy can be considered a pure strategy, and, with finitely many—in particular $O(1/\epsilon)$ —pure strategies, an anonymous game can be solved exhaustively in polynomial time in n , the number of players. The only problem is, why should the expected utilities before and after the quantization be close? Here we rely on a probabilistic lemma (Theorem 3.1) that may be of much more general interest: Given n Bernoulli random variables with probabilities p_1, \dots, p_n , there is a way to round the probabilities to multiples of $1/k$, for any k , so that the distribution of the sum of these n variables is affected only by an additive $O(1/\sqrt{k})$ in total variational distance (no dependence on n). This implies that the expected utilities of the quantized version are within an additive $\pm O(1/\sqrt{k})$ of the original ones, and an $O(n^{1/\epsilon^2})$ PTAS for two-strategy anonymous games is immediate. We feel that a more sophisticated proof of the same kind can establish a similar result for multinomial distributions, thus extending our PTAS to anonymous games with any fixed number of strategies.

1.1 Definitions and Notation

An *anonymous game* $G = (n, s, \{u_i^p\})$ consists of a set $[n] = \{1, \dots, n\}$ of $n \geq 2$ of players, a set $[s] = \{1, \dots, s\}$ of $s \geq 2$ strategies, and a set of ns utility functions, where u_i^p with $p \in [n]$ and $i \in [s]$ is the utility of player p when she plays strategy i , a function mapping the set of partitions $\Pi_{n-1}^s = \{(x_1, \dots, x_s) : x_i \in N_0 \text{ for all } i \in [s], \sum_{i=1}^s x_i = n - 1\}$ to the interval $[0, 1]$ ¹. Our working assumptions are that n is large and s is fixed; notice that, in this case, anonymous games are *succinctly representable* [23], in the sense that their representation requires specifying $O(n^{s+1})$ numbers, as opposed to the ns^n numbers required for general games (arguably, succinct games are the only multiplayer games that are computationally meaningful, see [23] for an extensive discussion of this point). For our approximate pure Nash equilibrium result we shall also be assuming that the utility functions are continuous, in the following sense: There is a real $\lambda > 0$, presumably very small, such that $|u_i^p(x) - u_i^p(y)| \leq \lambda \cdot \|x - y\|_1$ for every $p \in [n], i \in [s]$, and $x, y \in \Pi_{n-1}^s$. This continuity concept is similar to the “bounded jump” assumption of [8]. The convex hull of the set Π_{n-1}^s will be denoted by $\Delta_{n-1}^s = \{(x_1, \dots, x_s) : x_i \geq 0, i = 1, \dots, s, \sum_{i=1}^s x_i = n - 1\}$.

¹In the literature on Nash approximation utilities are usually normalized this way so that the approximation error is additive.

A *pure strategy profile* in such a game is a mapping S from $[n]$ to $[s]$. A pure strategy profile S is an ϵ -approximate pure Nash equilibrium, where $\epsilon > 0$, if, for all $p \in [n]$, $u_{S(p)}^p(x[S, p]) + \epsilon \geq u_i^p(x[S, p])$ for all $i \in [s]$, where $x[S, p] \in \Pi_{n-1}^s$ is the partition (x_1, \dots, x_s) such that x_i is the number of players $q \in [n] - \{p\}$ such that $S(q) = i$.

A *mixed strategy profile* is a set of n distributions $\delta_p, p \in [n]$, over $[s]$. A mixed strategy profile is an ϵ -approximate mixed Nash equilibrium if, for all $p \in [n]$ and $j \in [s]$, $E_{\delta_1, \dots, \delta_n} u_i^p(x) + \epsilon \geq E_{\delta_1, \dots, \delta_n} u_j^p(x)$ where, for the purposes of the expectation, i is drawn from $[s]$ according to δ_p and x is drawn from Π_{n-1}^s by drawing $n-1$ random samples from $[s]$ independently according to the distributions $\delta_q, q \neq p$ and forming the induced partition.

Anonymous games can be extended to ones in which there is also a finite number of *types* of players, and utilities depend on how each type is partitioned into strategies; all our algorithms, being exhaustive, can be easily generalized to this framework, with the number of types multiplying the exponent.

2 Approximate Pure Equilibria

In this section we prove the following result:

Theorem 2.1 *In any anonymous game with s strategies and Lipschitz constant λ there is an ϵ -approximate pure Nash equilibrium, where $\epsilon = O(s^2)\lambda$.*

Proof: We first define a function ϕ from Π_{n-1}^s to itself: For any $x \in \Pi_{n-1}^s$, $\phi(x)$ is defined to be $(y_1, \dots, y_s) \in \Pi_{n-1}^s$ such that, for all $i \in [s]$, y_i is the number of all those players p among $\{1, \dots, n-1\}$ (notice that player n is excluded) such that, for all $j < i$, $u_i^p[x] > u_j^p[x]$, and, for all $j > i$, $u_i^p[x] \geq u_j^p[x]$. In other words, $\phi(x)$ is the partition induced among the first $n-1$ players by their best response to x , where ties are broken lexicographically.

We next interpolate ϕ to obtain a continuous function $\hat{\phi}$ from Δ_{n-1}^s to itself as follows: For each $x \in \Delta_{n-1}^s$ let us break x into its integer and fractional parts $x = x^I + x^F$, where $x^I \in \Pi_{n-1}^s$ and $0 \leq x_i^F \leq 1$ for all $i = 1, \dots, s-1$. Let $c[x^I]$ be the *cell* of x^I , the set of all $x' \in \Pi_{n-1}^s$ such that, for all $i \leq s-1$, $x'_i = x_i^I$ or $x'_i = x_i^I + 1$. Then it is clear that x can be written as a convex combination of the elements of $c[x^I]$: $x = \sum_{x^j \in c[x^I]} \alpha_j x^j$. We define $\hat{\phi}(x)$ to be $\sum_{x^j \in c[x^I]} \alpha_j \phi(x^j)$.

It is possible to define the interpolation at each point $x \in \Delta_{n-1}^s$ in a consistent way so that the resulting $\hat{\phi}(x)$ is a continuous function from the compact set Δ_{n-1}^s to itself, and so, by Brouwer's Theorem, it must have

a fixed point, that is, a point $x^* \in \Delta_{n-1}^s$ such that $\hat{\phi}(x^*) = x^*$. That is,

$$\sum_{x^j \in c[x^*]^I} \alpha_j \phi(x^j) = x^* = \sum_{x^j \in c[x^*]^I} \alpha_j x^j. \quad (1)$$

By Carathéodory's lemma, equation (1) can be expressed as the sum of only s of the $\phi(x^j)$'s $x^* = \sum_{j=1}^s \gamma_j \phi(x^j)$, and it is easy to see that it can be rewritten as

$$x^* = \phi(x^1) + \sum_{j=2}^s \gamma_j (\phi(x^j) - \phi(x^1)), \quad (2)$$

for some $\gamma_j \geq 0$ with $\sum_j \gamma_j = 1$.

Recall that, in order to prove the theorem, we need to exhibit ϵ -approximate pure strategy profile. If x^* were an integer point, then we would be almost done (modulo the n -th player, of whom we take care last), and x^* itself (actually, the strategy profile suggested by the partition $\phi(x^*)$) would be essentially a pure Nash equilibrium, because of the equation $x^* = \phi(x^*)$. But in general x^* will be fractional, and the various $\phi(x^j)$'s will be very far from x^* (except that they happen to have x^* in their convex hull). Our plan is to show that x^1 (a vertex in the cell of x^*) is an approximate pure Nash equilibrium (again, considered as a pure strategy profile and forgetting for a moment the n -th player).

The term $\phi(x^1)$ in equation (2) can be seen as a pure strategy profile P : Each of the $n-1$ players chooses the strategy that is her/his best response to x^1 . Therefore, in this strategy profile everybody would be happy if everybody else played according to x^1 . The problem is, of course, that $\phi(x^1)$ can be very far from x^1 . We shall next use equation (2) to “move it” close to x^1 (more precisely, close to x^* which we know is $2s$ -close in L_1 distance to x^1) without changing the utilities much. Looking at one of the other terms of (2), $(\phi(x^j) - \phi(x^1))$, we can think of it as the act of switching certain players from their best response to x^1 to their best response to x^j . The crucial observation is that, *since x^1 and x^j are at most $2s$ apart in L_1 distance (they both belong to the same cell), the change in utility for the switching players would be at most $4s\lambda$.*

So, equation (2) suggests that a strategy profile close to x^* can be obtained from P by combining these $s-1$ flows, with little harm in utility for all players involved. The problem is how to combine them so that the right individual players are switched (the situation is akin to integer multicommodity flow). We write each flow $(\phi(x^j) - \phi(x^1))$ as the sum of s^2 terms of the form $f_{i,i'}^j$, signifying the number of individual players moved from strategy i to strategy i' . We know that, for each such nonzero flow, there is a set of players $S_{i,i'}^j \subseteq [n]$ which can be moved with only $4s\lambda$ loss in utility. The union

over j of the $(s - 1)$ sets $S_{i,i'}^j$ is denoted by $S_{i,i'}$ and the union over i' of the sets $S_{i,i'}$ by S_i . The following lemma can be proved by an application of Hall's Theorem.

Lemma 2.2 *There exist disjoint subsets $T_{i,i'} \subseteq S_{i,i'}$, $i' = 1, \dots, s$, such that, for all $i' = 1, \dots, s$, $|T_{i,i'}| = \lfloor \sum_{j=2}^s \gamma_j f_{i,i'}^j \rfloor$.*

Proof: Let us consider the bipartite graph with vertex set $S_i \sqcup [s]$ and an edge from a player $p \in S_i$ to a strategy $i' \in [s]$ if, for some j , $p \in S_{i,i'}^j$. To establish the result it is enough to show that there exists a generalized matching of players to pure strategies—in which every player is matched to at most one pure strategy, so that, for all i' , strategy i' is matched with at least $R_{i'} := \lfloor \sum_{j=2}^s \gamma_j f_{i,i'}^j \rfloor$ players.

By Hall's theorem, such a matching exists if every set of strategies $\mathcal{J} \subseteq [s]$ “knows” at least $\sum_{i' \in \mathcal{J}} R_{i'}$ players. Observe that, for all j , the family of sets $\{S_{i,i'}^j\}_{i'}$ are disjoint. Hence

$$\sum_{i' \in \mathcal{J}} f_{i,i'}^j \leq \left| \bigcup_{i' \in \mathcal{J}} S_{i,i'}^j \right| \leq \left| \bigcup_{i' \in \mathcal{J}} S_{i,i'} \right| = |\Gamma(\mathcal{J})|,$$

where $\Gamma(\mathcal{J})$ represents the neighborhood in S_i of the pure strategies of the set \mathcal{J} . From the above equation it follows that

$$\begin{aligned} \sum_j \gamma_j \sum_{i' \in \mathcal{J}} f_{i,i'}^j &\leq \sum_j \gamma_j |\Gamma(\mathcal{J})| \leq |\Gamma(\mathcal{J})| \\ \Rightarrow \sum_{i' \in \mathcal{J}} \sum_j \gamma_j f_{i,i'}^j &\leq |\Gamma(\mathcal{J})| \Rightarrow \sum_{i' \in \mathcal{J}} R_{i'} \leq |\Gamma(\mathcal{J})|, \end{aligned}$$

which completes the proof. ■

Thus, by moving the players in $T_{i,i'}$ from i to i' , for all pairs of i, i' , we obtain from strategy profile P a new strategy profile \tilde{P} in which each player's strategy is within $4s\lambda$ from their response to x^1 , and such that the corresponding partition \tilde{x}^* is, by equation (2) and the roundings in the lemma, at most $2s^2$ away from x^* , and hence at most $2s$ more away from x^1 ; let's call the distance bound $D = 2s^2 + 2s$. Since, for all players—except for the last of course, \tilde{P} is an $4s\lambda$ -approximate response to x^1 and \tilde{P} is within distance D from x^1 , it follows that \tilde{P} is a $4(D+s)\lambda$ -approximate best response to itself.

Finally, we turn to player n . Adding the best response of player n to \tilde{P} , and subtracting what player i plays in \tilde{P} , we get a profile that is 2 away, in L_1 distance, from \tilde{P} , thus making \tilde{P} a $4(D+s+1)\lambda$ -approximate Nash equilibrium and completing the proof. ■

Since Π_{n-1}^s has $O(n^s)$ points, and this is the length of the input, the algorithmic implication is immediate:

Corollary 2.3 *In any anonymous game, an ϵ -approximate pure Nash equilibrium, where ϵ is as in Theorem 2.1, can be found in linear time.*

3 Approximate Mixed Nash Equilibria

3.1 A Probabilistic Lemma

We start by a definition. The *total variation distance* between two distributions \mathbb{P} and \mathbb{Q} supported on a finite set \mathcal{A} is

$$\|\mathbb{P} - \mathbb{Q}\| = \frac{1}{2} \sum_{\alpha \in \mathcal{A}} |\mathbb{P}(\alpha) - \mathbb{Q}(\alpha)|.$$

Theorem 3.1 *Let $\{p_i\}_{i=1}^n$ be arbitrary probabilities, $p_i \in [0, 1]$, for $i = 1, \dots, n$, and let $\{X_i\}_{i=1}^n$ be independent indicator random variables, such that X_i has expectation $\mathcal{E}[X_i] = p_i$, and let k be a positive integer. Then there exists another set of probabilities $\{q_i\}_{i=1}^n$, $q_i \in [0, 1]$, $i = 1, \dots, n$, which satisfy the following properties:*

1. $\|q_i - p_i\| = O(1/k)$, for all $i = 1, \dots, n$
2. q_i is an integer multiple of $\frac{1}{k}$, for all $i = 1, \dots, n$
3. if $\{Y_i\}_{i=1}^n$ are independent indicator random variables such that Y_i has expectation $\mathcal{E}[Y_i] = q_i$, then,

$$\left\| \sum_i X_i - \sum_i Y_i \right\| = O(k^{-1/2}).$$

and, moreover, for all $j = 1, \dots, n$,

$$\left\| \sum_{i \neq j} X_i - \sum_{i \neq j} Y_i \right\| = O(k^{-1/2}).$$

From this, the main result of this section follows:

Corollary 3.2 *There is a PTAS for the mixed Nash equilibrium problem for two-strategy anonymous games.*

Proof: Let (p_1, \dots, p_n) be a mixed Nash equilibrium of the game. We claim that (q_1, \dots, q_n) , where the q_i 's are the multiples of $1/k$ specified by Theorem 3.1, constitute a $O(1/\sqrt{k})$ -approximate mixed Nash equilibrium. Indeed, for every player $i \in [n]$ and every strategy $m \in \{1, 2\}$ for that player let us track the change in the expected utility of the player when the distribution over Π_{n-1}^2 defined by the $\{p_j\}_{j \neq i}$ is replaced by the distribution defined by the $\{q_j\}_{j \neq i}$. It is not hard to see that the absolute change is bounded by the total variation distance between the distributions of the $\sum_{j \neq i} X_j$

and the $\sum_{j \neq i} Y_j$ ² where X_j, Y_j are indicators corresponding to whether player j plays strategy 2 in the distribution defined by the p_i 's and the q_i 's respectively, i.e. $\mathcal{E}[X_j] = p_j$ and $\mathcal{E}[Y_j] = q_j$. Hence, the change in utility is at most $O(1/\sqrt{k})$, which implies that the q_i 's constitute an $O(1/\sqrt{k})$ -approximate Nash equilibrium of the game, modulo the following observation: with a trivial modification in the proof of Theorem 3.1 we can ensure that, when switching from p_i 's to q_i 's, for every i , the support of q_i is a subset of the support of p_i .

To compute a quantized approximate Nash equilibrium of the original game, we proceed to define a related $(k+1)$ -strategy game, where $k = O(\frac{1}{\epsilon^2})$, and treat the problem as a pure Nash equilibrium problem. It is not hard to see that the latter is efficiently solvable if the number of strategies is a constant. The new game is defined as follows: the i -th pure strategy, $i = 0, \dots, k$, corresponds to a player in the original game playing strategy 2 with probability $\frac{i}{k}$. Naturally, the payoffs resulting from a pure strategy profile in the new game are defined to be equal to the corresponding payoffs in the original game, by the translation of the pure strategy profile of the former into a mixed strategy profile of the latter. In particular, for any player p , we can compute its payoff given any strategy $i \in \{0, \dots, k\}$ for that player and any partition $x \in \prod_{n=1}^{k+1}$ of the other players into $k+1$ strategies, in time $n^{O(1/\epsilon^2)}$ overall, by a straightforward dynamic programming algorithm, see for example [24]. The remaining details are omitted. ■

Remark: Note that it is crucial for the proof of Corollary 3.2 that the bound on the total variation distance between the $\sum_i X_i$ and the $\sum_i Y_i$ in the statement of the Theorem 3.1 does not depend on the number n of random variables which are being rounded, but only on the accuracy $\frac{1}{k}$ of the rounding. Because of this requirement, several simple methods of rounding are easily seen to fail:

- *Rounding to the Closest Multiple of $1/k$:* An easy counterexample for this method arises when $p_i := \frac{1}{n}$, for all i . In this case, the trivial rounding would make $q_i := 0$, for all i , and the total variation distance between the $\sum_i X_i$ and the $\sum_i Y_i$ would become arbitrarily close to $1 - \frac{1}{e}$, as n goes to infinity.
- *Randomized Rounding:* An argument employing the probabilistic method could start by independently rounding each p_i to some random q_i which is an integer multiple of $\frac{1}{k}$ in such a way that $\mathcal{E}[q_i] = p_i$. This seems promising since, by independence, for any $\ell = 0, \dots, n$, the random variable $\Pr[\sum_i Y_i = \ell]$, which is a function of the q_i 's,

has the correct expectation, i.e. $\mathcal{E}[\Pr[\sum_i Y_i = \ell]] = \Pr[\sum_i X_i = \ell]$. The trouble is that the expectation of the random variable $\Pr[\sum_i Y_i = \ell]$ is very small: less than 1 for all ℓ and, in fact, in the order of $O(1/n)$ for many terms. Moreover, the function itself comprises of sums of products on the random variables q_i , in fact exponentially many terms for some values of ℓ . Concentration seems to require k which scales polynomially in n .

Proof Technique: We follow instead a completely different approach which aims at directly approximating the distribution of the $\sum_i X_i$. The intuition is the following: The distribution of the $\sum_i X_i$ should be close in total variation distance to a Poisson distribution of the same mean $\sum_i p_i$. Hence, it seems that, if we define q_i 's—which are multiples of $\frac{1}{k}$ —in such a way that the means $\sum_i p_i$ and $\sum_i q_i$ are close, then the distribution of the $\sum_i Y_i$ should be close in total variation distance to the same Poisson distribution and hence to the distribution of the $\sum_i X_i$ by triangle inequality.

There are several complications, of course, the main one being that the distribution of the $\sum_i X_i$ can be well approximated by a Poisson distribution of the same mean only when the p_i 's are relatively small. When the p_i 's take arbitrary values in $[0, 1]$ and n scales, the Poisson distribution can be very far from the distribution of the $\sum_i X_i$. In fact, we wouldn't expect that the Poisson distribution can approximate the distribution of arbitrary sums of indicators since its mean and variance are the same. To counter this we resort to a special kind of distributions, called *translated Poisson distributions*, which are Poisson distributions appropriately shifted on their domain. An arbitrary sum of indicators can be now approximated as follows: a Poisson distribution is defined with mean — and, hence, variance — equal to the variance of the sum of the indicators; then the distribution is appropriately shifted on its domain so that its new mean coincides with the mean of the sum of the indicators being approximated.

The translated Poisson approximation will outperform the Poisson approximation for intermediate values of the p_i 's, while the Poisson approximation will remain better near the boundaries, i.e. for values of p_i close to 0 or 1. Even for the intermediate region of values for the p_i 's, the translated Poisson approximation is not sufficient since it only succeeds when the number of the indicators being summed over is relatively large, compared to the minimum expectation. A different argument is required when this is not the case. Our bounding technique has to interleave these considerations in a very delicate fashion to achieve the approximation result. At a high level, we treat separately the X_i 's with small, medium or large expectation; in particular, for some $\alpha \in (0, 1)$

²Recall that all utilities have been normalized to take values in $[0, 1]$.

to be fixed later, we define the following subintervals of $[0, 1]$:

1. $\mathcal{L}(k) := \left[0, \frac{\lfloor k^\alpha \rfloor}{k}\right)$: interval of *small expectations*;
2. $\mathcal{M}_1(k) := \left[\frac{\lfloor k^\alpha \rfloor}{k}, \frac{k/2}{k}\right)$: first interval of *medium expectations*;
3. $\mathcal{M}_2(k) := \left[\frac{k/2}{k}, 1 - \frac{\lfloor k^\alpha \rfloor}{k}\right)$: second interval of *medium expectations*;
4. $\mathcal{H}(k) := \left[1 - \frac{\lfloor k^\alpha \rfloor}{k}, 1\right]$: interval of *high expectations*.

Denoting $\mathcal{L}^*(k) := \{i \mid \mathcal{E}[X_i] \in \mathcal{L}(k)\}$, we establish (Lemma 3.9) that

$$\left\| \sum_{i \in \mathcal{L}^*(k)} X_i - \sum_{i \in \mathcal{L}^*(k)} Y_i \right\| = O(k^{-1/2})$$

and similarly for $\mathcal{M}_1(k)$ (Lemma 3.12). Symmetric arguments (setting $X'_i = 1 - X_i$ and $Y'_i = 1 - Y_i$) imply the same bounds for the intervals $\mathcal{M}_2(k)$ and $\mathcal{H}(k)$. Therefore, an application of the coupling lemma implies that

$$\left\| \sum_i X_i - \sum_i Y_i \right\| = O(k^{-1/2}),$$

which concludes the proof. The details of the proof are postponed to Section 3.3. The proof for the partial sums $\sum_{i \neq j} X_i$ and $\sum_{i \neq j} Y_i$ follows easily from the analysis of Section 3.3 and its details are skipped for this extended abstract. The next section provides the required background on Poisson approximations.

3.2 Poisson Approximations

The following theorem is classical in the theory of Poisson approximations.

Theorem 3.3 ([2]) *Let J_1, \dots, J_n be a sequence of independent random indicators with $\mathcal{E}[J_i] = p_i$. Then*

$$\left\| \sum_{i=1}^n J_i - \text{Poisson} \left(\sum_{i=1}^n p_i \right) \right\| \leq \frac{\sum_{i=1}^n p_i^2}{\sum_{i=1}^n p_i}.$$

As discussed in the previous section, the above bound is sharp when the indicators have small expectations, but loose when the indicators are arbitrary. The following approximation bound becomes sharp when the previous is not. But first let us formally define the translated Poisson distribution.

Definition 3.4 ([25]) We say that an integer random variable Y has a *translated Poisson distribution* with parameters μ and σ^2 and write

$$\mathcal{L}(Y) = TP(\mu, \sigma^2)$$

if $\mathcal{L}(Y - \lfloor \mu - \sigma^2 \rfloor) = \text{Poisson}(\sigma^2 + \{\mu - \sigma^2\})$, where $\{\mu - \sigma^2\}$ represents the fractional part of $\mu - \sigma^2$.

Theorem 3.5 provides an approximation result for the translated Poisson distribution using Stein's method.

Theorem 3.5 ([25]) *Let J_1, \dots, J_n be a sequence of independent random indicators with $\mathcal{E}[J_i] = p_i$. Then*

$$\left\| \sum_{i=1}^n J_i - TP(\mu, \sigma^2) \right\| \leq \frac{\sqrt{\sum_{i=1}^n p_i^3 (1 - p_i)} + 2}{\sum_{i=1}^n p_i (1 - p_i)},$$

where $\mu = \sum_{i=1}^n p_i$ and $\sigma^2 = \sum_{i=1}^n p_i (1 - p_i)$.

Lemmas 3.6 and 3.7 provide respectively bounds for the total variation distance between two Poisson distributions and two translated Poisson distributions with different parameters. The proof of 3.6 is postponed to the appendix, while the proof of 3.7 is provided in [3].

Lemma 3.6 *Let $\lambda_1, \lambda_2 \in \mathbb{R}_+ \setminus \{0\}$. Then*

$$\begin{aligned} & \left| \text{Poisson}(\lambda_1) - \text{Poisson}(\lambda_2) \right| \\ & \leq e^{|\lambda_1 - \lambda_2|} - e^{-|\lambda_1 - \lambda_2|}. \end{aligned}$$

Lemma 3.7 ([3]) *Let $\mu_1, \mu_2 \in \mathbb{R}$ and $\sigma_1^2, \sigma_2^2 \in \mathbb{R}_+ \setminus \{0\}$ be such that $\lfloor \mu_1 - \sigma_1^2 \rfloor \leq \lfloor \mu_2 - \sigma_2^2 \rfloor$. Then*

$$\begin{aligned} & \left| TP(\mu_1, \sigma_1^2) - TP(\mu_2, \sigma_2^2) \right| \\ & \leq \frac{|\mu_1 - \mu_2|}{\sigma_1} + \frac{|\sigma_1^2 - \sigma_2^2| + 1}{\sigma_1^2}. \end{aligned}$$

3.3 Proof Theorem 3.1

In this section we complete the proof of Theorem 3.1. As argued above, it is enough to round the random variables $\{X_i\}_{i \in \mathcal{L}^*(k)}$ into random variables $\{Y_i\}_{i \in \mathcal{L}^*(k)}$ so that the total variation distance between the random variables $\mathcal{X}_{\mathcal{L}} := \sum_{i \in \mathcal{L}^*(k)} X_i$ and $\mathcal{Y}_{\mathcal{L}} := \sum_{i \in \mathcal{L}^*(k)} Y_i$ is small and similarly for the subinterval $\mathcal{M}_1(k)$.

Our rounding will have different objective in the two regions. When rounding the X_i 's with $i \in \mathcal{L}^*(k)$ we aim to approximate the mean of $\mathcal{X}_{\mathcal{L}}$ as tightly as possible. On the other hand, when rounding the X_i 's with $i \in \mathcal{M}_1^*(k)$, we give up on approximating the mean very tightly in order to also approximate well the variance. The details of the rounding follow.

Some notation first: Let us partition the interval $[0, 1/2]$ into $\lceil k/2 \rceil$ subintervals $I_0, I_1, \dots, I_{\lfloor k/2 \rfloor}$ where

$$I_0 = \left[0, \frac{1}{k}\right), I_1 = \left[\frac{1}{k}, \frac{2}{k}\right), \dots, I_{\lfloor k/2 \rfloor} = \left[\frac{\lfloor k/2 \rfloor}{k}, \frac{1}{2}\right].$$

The intervals $I_0, \dots, I_{\lfloor k/2 \rfloor}$ define the partition of $\mathcal{L}^*(k) \cup \mathcal{M}_1^*(k)$ into the subsets $I_0^*, I_1^*, \dots, I_{\lfloor k/2 \rfloor}^*$, where

$$I_j^* = \{i \mid \mathcal{E}[X_i] \in I_j\}, j = 0, 1, \dots, \lfloor k/2 \rfloor.$$

For all $j \in \{0, \dots, \lfloor k/2 \rfloor\}$ with $I_j^* \neq \emptyset$, let $I_j^* = \{j_1, j_2, \dots, j_{n_j}\}$ and, for all $i \in \{1, \dots, n_j\}$, let

$$p_i^j := \mathcal{E}[X_{j_i}] \text{ and } \delta_i^j := p_i^j - \frac{j}{k}.$$

We proceed to define the “rounding” of the X_i ’s into the Y_i ’s in the intervals $\mathcal{L}(k)$ and $\mathcal{M}_1(k)$ separately.

Interval $\mathcal{L}(k) := [0, \frac{\lfloor k^\alpha \rfloor}{k}]$ of small expectations.

Observe first that $\mathcal{L}(k) \equiv I_0 \cup \dots \cup I_{\lfloor k^\alpha \rfloor - 1}$ and define the corresponding subset of the indices $\mathcal{L}^*(k) := I_0^* \cup \dots \cup I_{\lfloor k^\alpha \rfloor - 1}^*$. We define the Y_i , $i \in \mathcal{L}^*(k)$, via the following iterative procedure. Our ultimate goal is to round the X_i ’s into Y_i ’s appropriately so that the sum of the expectations of the X_i ’s and of the Y_i ’s are as close as possible. The rounding procedure is as follows.

- i. $\epsilon_0 := 0$;
- ii. for $j := 0$ to $\lfloor k^\alpha \rfloor - 1$
 - (a) $S_j := \epsilon_j + \sum_{i=1}^{n_j} \delta_i^j$;
 - (b) $m_j := \left\lfloor \frac{S_j}{k^{-1}} \right\rfloor$; $\epsilon_{j+1} := S_j - m_j \cdot \frac{1}{k}$;
{assertion: $m_j \leq n_j$ - see justification next}
 - (c) set $q_i^j := \frac{j+1}{k}$ for $i = 1, \dots, m_j$ and $q_i^j := \frac{j}{k}$ for $i = m_j + 1, \dots, n_j$;
 - (d) for all $i \in \{1, \dots, n_j\}$, let Y_{j_i} be a $\{0, 1\}$ -random variable with expectation q_i^j ;
- iii. Suppose that the random variables Y_i , $i \in \mathcal{L}^*(k)$, are mutually independent.

It is easy to see that, for all $j \in \{0, \dots, \lfloor k^\alpha \rfloor - 1\}$, $\epsilon_j < \frac{1}{k}$; this follows immediately from the description of the procedure, in particular Steps i and ii(b). This further implies that $m_j \leq n_j$, for all j , since at Step ii(b) we have

$$\frac{S_j}{k^{-1}} = \frac{\epsilon_j + \sum_{i=1}^{n_j} \delta_i^j}{k^{-1}} < \frac{k^{-1} + n_j k^{-1}}{k^{-1}} = 1 + n_j.$$

Hence, the assertion following Step ii(b) is satisfied. Finally, note that, for all j ,

$$\begin{aligned} \sum_{i=1}^{n_j} q_i^j &= m_j \frac{j+1}{k} + (n_j - m_j) \frac{j}{k} \\ &= n_j \frac{j}{k} + m_j \frac{1}{k} = n_j \frac{j}{k} + S_j - \epsilon_{j+1} \\ &= n_j \frac{j}{k} + \sum_{i=1}^{n_j} \delta_i^j + \epsilon_j - \epsilon_{j+1} = \sum_{i=1}^{n_j} p_i^j + \epsilon_j - \epsilon_{j+1}. \end{aligned}$$

Therefore,

$$\sum_{j=0}^{\lfloor k^\alpha \rfloor - 1} \sum_{i=1}^{n_j} q_i^j = \sum_{j=0}^{\lfloor k^\alpha \rfloor - 1} \sum_{i=1}^{n_j} p_i^j + \epsilon_0 - \epsilon_{\lfloor k^\alpha \rfloor},$$

which implies

Lemma 3.8 $|\sum_{i \in \mathcal{L}^*(k)} \mathcal{E}[Y_i] - \sum_{i \in \mathcal{L}^*(k)} \mathcal{E}[X_i]| \leq \frac{1}{k}.$

The following lemma characterizes the total variation distance between $\sum_{i \in \mathcal{L}^*(k)} X_i$ and $\sum_{i \in \mathcal{L}^*(k)} Y_i$.

Lemma 3.9 $\left\| \sum_{i \in \mathcal{L}^*(k)} X_i - \sum_{i \in \mathcal{L}^*(k)} Y_i \right\| \leq \frac{3}{k^{1-\alpha}}.$

Proof: By lemma 3.3 we have that

$$\left\| \sum_{i \in \mathcal{L}^*(k)} X_i - \text{Poisson} \left(\sum_{i \in \mathcal{L}^*(k)} p_i \right) \right\| \leq \frac{\sum_{i \in \mathcal{L}^*(k)} p_i^2}{\sum_{i \in \mathcal{L}^*(k)} p_i}$$

and

$$\left\| \sum_{i \in \mathcal{L}^*(k)} Y_i - \text{Poisson} \left(\sum_{i \in \mathcal{L}^*(k)} q_i \right) \right\| \leq \frac{\sum_{i \in \mathcal{L}^*(k)} q_i^2}{\sum_{i \in \mathcal{L}^*(k)} q_i},$$

where $p_i := \mathcal{E}[X_i]$ and $q_i := \mathcal{E}[Y_i]$ for all i . The following lemma is proven in the appendix.

Lemma 3.10 For any $u > 0$ and any set $\{p_i\}_{i \in \mathcal{I}}$, where $p_i \in [0, u]$, for all $i \in \mathcal{I}$,

$$\frac{\sum_{i \in \mathcal{I}} p_i^2}{\sum_{i \in \mathcal{I}} p_i} \leq u.$$

Using Lemmas 3.10, 3.8 and 3.6 and the triangle inequality we get that

$$\begin{aligned} \left\| \sum_{i \in \mathcal{L}^*(k)} X_i - \sum_{i \in \mathcal{L}^*(k)} Y_i \right\| &\leq \frac{2}{k^{1-\alpha}} + (e^{\frac{1}{k}} - e^{-\frac{1}{k}}) \\ &\leq \frac{3}{k^{1-\alpha}}, \end{aligned}$$

where we used the fact that $e^{\frac{1}{k}} - e^{-\frac{1}{k}} \leq \frac{3}{k} \leq \frac{1}{k^{1-\alpha}}$, for sufficiently large k . ■

Interval $\mathcal{M}_1(k) := \left[\lfloor \frac{k^\alpha}{k} \rfloor, \frac{k/2}{k} \right)$: **medium expectations.**

Observe first that $\mathcal{M}_1(k) \equiv I_{\lfloor k^\alpha \rfloor} \cup \dots \cup I_{\lfloor k/2 \rfloor}$ and define the corresponding subset of the indices $\mathcal{M}_1^*(k) := I_{\lfloor k^\alpha \rfloor}^* \cup \dots \cup I_{\lfloor k/2 \rfloor}^*$. We define the $Y_i, i \in \mathcal{M}_1^*(k)$, via the following procedure which is slightly different than the one we used for the set of indices $\mathcal{L}^*(k)$. Our goal here is to approximate well both the mean and the variance of the sum $\sum_{i \in \mathcal{M}_1^*(k)} X_i$. In fact, we will give up on approximating the mean as tightly as possible, which we did above, in order to achieve a good approximation of the variance. The rounding procedure is as follows.

for $j := \lfloor k^\alpha \rfloor$ to $\lfloor \frac{k}{2} \rfloor$

- (a) $S_j := \sum_{i=1}^{n_j} \delta_i^j$;
- (b) $m_j := \left\lfloor \frac{S_j}{k^{-1}} \right\rfloor$;
- (c) set $q_i^j := \frac{j+1}{k}$ for $i = 1, \dots, m_j$ and $q_i^j := \frac{j}{k}$ for $i = m_j + 1, \dots, n_j$;
- (d) for all $i \in \{1, \dots, n_j\}$, let Y_i be a $\{0, 1\}$ -random variable with expectation q_i^j ;

Suppose that the random variables $Y_i, i \in \mathcal{M}_1^*(k)$, are mutually independent.

Lemma 3.11 characterizes the quality of the rounding procedure in terms of mean and variance. Defining $\zeta_j := \sum_{i \in I_j^*} \mathcal{E}[X_i] - \sum_{i \in I_j^*} \mathcal{E}[Y_i]$, we have

Lemma 3.11 For all $j \in \{\lfloor k^\alpha \rfloor, \dots, \lfloor \frac{k}{2} \rfloor\}$

- (a) $\zeta_j = \sum_{i=1}^{n_j} \delta_i^j - m_j \frac{1}{k}$.
- (b) $0 \leq \zeta_j \leq \frac{1}{k}$.
- (c) $\sum_{i \in I_j^*} \text{Var}[X_i] = n_j \frac{j}{k} \left(1 - \frac{j}{k}\right) + \left(1 - \frac{2j}{k}\right) \sum_{i=1}^{n_j} \delta_i^j - \sum_{i=1}^{n_j} (\delta_i^j)^2$
- (d) $\sum_{i \in I_j^*} \text{Var}[Y_i] = n_j \frac{j}{k} \left(1 - \frac{j}{k}\right) + m_j \frac{1}{k} \left(1 - \frac{2j+1}{k}\right)$
- (e) $\sum_{i \in I_j^*} \text{Var}[X_i] - \sum_{i \in I_j^*} \text{Var}[Y_i] = \left(1 - \frac{2j}{k}\right) \zeta_j + \left(m_j \frac{1}{k^2} - \sum_{i=1}^{n_j} (\delta_i^j)^2\right)$.

The following lemma bounds the total variation distance between the random variables $\sum_{i \in \mathcal{M}_1^*(k)} X_i$ and $\sum_{i \in \mathcal{M}_1^*(k)} Y_i$.

Lemma 3.12 $\left\| \sum_{i \in \mathcal{M}_1^*(k)} X_i - \sum_{i \in \mathcal{M}_1^*(k)} Y_i \right\| \leq O\left(k^{-\frac{\alpha+\beta-1}{2}}\right) + O(k^{-\alpha}) + O(k^{-\frac{1}{2}}) + O(k^{-(1-\beta)})$.

Proof: We distinguish two cases for the size of $\mathcal{M}_1^*(k)$. For some $\beta \in (0, 1)$ such that $\alpha + \beta > 1$, let us distinguish two possibilities for the size of $|\mathcal{M}_1^*(k)|$:

- a. $|\mathcal{M}_1^*(k)| \leq k^\beta$
- b. $|\mathcal{M}_1^*(k)| > k^\beta$

Let us treat each interval separately in the following lemmas.

Lemma 3.13 If $|\mathcal{M}_1^*(k)| \leq k^\beta$ then

$$\left\| \sum_{i \in \mathcal{M}_1^*(k)} X_i - \sum_{i \in \mathcal{M}_1^*(k)} Y_i \right\| \leq \frac{1}{k^{1-\beta}}.$$

Proof: The proof follows from the coupling lemma and an easy coupling argument. The details are postponed to the appendix. ■

Lemma 3.14 If $|\mathcal{M}_1^*(k)| > k^\beta$ then

$$\left\| \sum_{i \in \mathcal{M}_1^*(k)} X_i - \sum_{i \in \mathcal{M}_1^*(k)} Y_i \right\| \leq O\left(k^{-\frac{\alpha+\beta-1}{2}}\right) + O(k^{-\alpha}) + O(k^{-\frac{1}{2}}).$$

Proof: By lemma 3.5 we have that

$$\left\| \sum_{i \in \mathcal{M}_1^*(k)} X_i - TP(\mu_1, \sigma_1^2) \right\| \leq \frac{\sqrt{\sum_{i \in \mathcal{M}_1^*(k)} p_i^3(1-p_i) + 2}}{\sum_{i \in \mathcal{M}_1^*(k)} p_i(1-p_i)}$$

and

$$\left\| \sum_{i \in \mathcal{M}_1^*(k)} Y_i - TP(\mu_2, \sigma_2^2) \right\| \leq \frac{\sqrt{\sum_{i \in \mathcal{M}_1^*(k)} q_i^3(1-q_i) + 2}}{\sum_{i \in \mathcal{M}_1^*(k)} q_i(1-q_i)}$$

where $\mu_1 = \sum_{i \in \mathcal{M}_1^*(k)} p_i$, $\mu_2 = \sum_{i \in \mathcal{M}_1^*(k)} q_i$, $\sigma_1^2 = \sum_{i \in \mathcal{M}_1^*(k)} p_i(1-p_i)$, $\sigma_2^2 = \sum_{i \in \mathcal{M}_1^*(k)} q_i(1-q_i)$ and $p_i := \mathcal{E}[X_i]$, $q_i := \mathcal{E}[Y_i]$ for all i . The following lemma is proven in the appendix.

Lemma 3.15 For any $u \in (0, \frac{1}{2})$ and any set $\{p_i\}_{i \in \mathcal{I}}$, where $p_i \in [u, \frac{1}{2}]$, for all $i \in \mathcal{I}$,

$$\frac{\sqrt{\sum_{i \in \mathcal{I}} p_i^3(1-p_i)}}{\sum_{i \in \mathcal{I}} p_i(1-p_i)} \leq \frac{(1+2u+4u^2-8u^3)}{\sqrt{16|\mathcal{I}|u(1-u-4u^2+4u^3)}}$$

Applying the above lemma with $u = \frac{\lfloor k^\alpha \rfloor}{k}$ and $\mathcal{I} = \mathcal{M}_1^*(k)$, where recall $|\mathcal{M}_1^*(k)| > k^\beta$ the above bound becomes

$$\frac{(1 + 2u + 4u^2 - 8u^3)}{\sqrt{16|\mathcal{I}|u(1 - u - 4u^2 + 4u^3)}} = O\left(k^{-\frac{\alpha+\beta-1}{2}}\right),$$

which implies

$$\left\| \sum_{i \in \mathcal{M}_1^*(k)} X_i - TP(\mu_1, \sigma_1^2) \right\| = O\left(k^{-\frac{\alpha+\beta-1}{2}}\right) \quad (1)$$

and

$$\left\| \sum_{i \in \mathcal{M}_1^*(k)} Y_i - TP(\mu_2, \sigma_2^2) \right\| = O\left(k^{-\frac{\alpha+\beta-1}{2}}\right), \quad (2)$$

where we used that, for any set of values $\{p_i \in \mathcal{M}_1(k)\}_{i \in \mathcal{M}_1^*(k)}$,

$$\frac{2}{\sum_{i \in \mathcal{M}_1^*(k)} p_i(1 - p_i)} \leq \frac{2}{|\mathcal{M}_1^*(k)| \frac{\lfloor k^\alpha \rfloor}{k} (1 - \frac{\lfloor k^\alpha \rfloor}{k})} = O\left(k^{-(\alpha+\beta-1)}\right)$$

and similarly for any set of values $\{q_i \in \mathcal{M}_1(k)\}_{i \in \mathcal{M}_1^*(k)}$.

All that remains to do is to bound the total variation distance between the distributions $TP(\mu_1, \sigma_1^2)$ and $TP(\mu_2, \sigma_2^2)$ for the parameters $\mu_1, \sigma_1^2, \mu_2, \sigma_2^2$ specified above. The following claim is proved in the appendix.

Claim 3.16 *For the parameters specified above*

$$\|TP(\mu_1, \sigma_1^2) - TP(\mu_2, \sigma_2^2)\| \leq O(k^{-\alpha}) + O(k^{-1/2}).$$

Claim 3.16 along with Equations (1) and (2) and the triangle inequality imply that

$$\left\| \sum_{i \in \mathcal{M}_1^*(k)} X_i - \sum_{i \in \mathcal{M}_1^*(k)} Y_i \right\| = O\left(k^{-\frac{\alpha+\beta-1}{2}}\right) + O(k^{-\alpha}) + O(k^{-1/2}).$$

■

■

Putting Everything Together.

Suppose that the random variables $\{Y_i\}_i$ defined above are mutually independent. It follows that

$$\left\| \sum_i X_i - \sum_i Y_i \right\| = O\left(k^{-(1-\alpha)}\right) + O\left(k^{-\frac{\alpha+\beta-1}{2}}\right) + O(k^{-\alpha}) + O(k^{-\frac{1}{2}}) + O(k^{-(1-\beta)}).$$

Setting $\alpha = \beta = \frac{3}{4}$ we get a total variation distance of $O(k^{-1/4})$. A more delicate argument establishes an exponent of $-\frac{1}{2}$.

4 Open Problems

Can our PTAS be extended to arbitrary fixed number of strategies? We believe so. A more sophisticated technique would subdivide, instead of the interval $[0, 1]$ as our proof did, the $(s-1)$ -dimensional simplex into domains in which multinomial (instead of binomial) distributions would be approximated in different ways, possibly using the techniques of [26]. This way of extending our result already seems to work for $s = 3$, and we are hopeful that it will work for general fixed s .

Can the quadratic, in s , approximation bound of our pure Nash equilibrium algorithm be improved to linear? We believe so, and we conjecture that $s\lambda$ is a lower bound.

We hope that the ways of thinking about anonymous games introduced in this paper will eventually lead to algorithms for the practical solution of this important class of games. Moreover, a technique involving probability rounding similar to the one used here yields a quasi-polynomial time approximation scheme for finding a Nash equilibrium in general normal form games with a fixed number of strategies, as well as for large classes of graphical games of this sort (work in progress). Improving this to polynomial is another important open problem.

Acknowledgment: We want to thank Uri Feige for a helpful discussion.

References

- [1] T. Abbott, D. Kane, P. Valiant. On the Complexity of Two-Player Win-Lose Games. *FOCS*, 2005.
- [2] A. D. Barbour, L. Holst and S. Janson. Poisson Approximation. *Oxford University Press*, New York, 1992.
- [3] A. D. Barbour and T. Lindvall. Translated Poisson Approximation for Markov Chains. *Journal of Theoretical Probability*, 19(3), July 2006.

- [4] M. Blonski. Characterization of Equilibria in Large Anonymous Games. University of Mannheim, 2000.
- [5] M. Blonski. Equilibrium Characterization in Large Anonymous Games. University of Mannheim, 2001.
- [6] X. Chen and X. Deng. Settling the Complexity of Two-Player Nash Equilibrium. *FOCS*, 2006.
- [7] X. Chen, X. Deng, and S.-H. Teng. Computing Nash Equilibria: Approximation and Smoothed Complexity. *FOCS*, 2006.
- [8] S. Chien and A. Sinclair. Convergence to Approximate Nash Equilibria in Congestion Games. *SODA*, 2007.
- [9] C. Daskalakis, P. Goldberg, and C. Papadimitriou. The Complexity of Computing a Nash Equilibrium. *STOC*, 2006.
- [10] C. Daskalakis, A. Mehta, and C. Papadimitriou. A Note on Approximate Nash Equilibria. *WINE*, 2006.
- [11] C. Daskalakis, A. Mehta, and C. Papadimitriou. Progress in Approximate Nash Equilibria. *EC*, 2007.
- [12] A. Fabrikant, C.H. Papadimitriou and K. Talwar. The Complexity of Pure Nash Equilibria. *STOC*, 2004.
- [13] T. Feder, H. Nazerzadeh, and A. Saberi. Approximating Nash Equilibria Using Small-Support Strategies. *EC*, 2007.
- [14] D. Gale, H. W. Kuhn, and A. W. Tucker. On Symmetric Games. In *H. W. Kuhn and A. W. Tucker, editors, Contributions to the Theory of Games*, 1:81–87, Princeton University Press, 1950.
- [15] P. Goldberg and C. Papadimitriou. Reducibility Among Equilibrium Problems. *STOC*, 2006.
- [16] E. Kalai. Partially-Specified Large Games. *WINE*, 2005.
- [17] R. Kannan and T. Theobald. Games of Fixed Rank: A Hierarchy of Bimatrix Games. *SODA*, 2007.
- [18] S. C. Kontogiannis, P. N. Panagopoulou, and P. G. Spirakis. Polynomial Algorithms for Approximating Nash Equilibria of Bimatrix Games. *WINE*, 2006.
- [19] R. Lipton, E. Markakis, and A. Mehta. Playing Large Games Using Simple Strategies. *Electronic Commerce*, 2003.
- [20] I. Milchtaich. Congestion Games with Player-Specific Payoff Functions. *Games and Economic Behavior*, 13:111–124.
- [21] J. Nash. Noncooperative Games. *Annals of Mathematics*, 54:289–295, 1951.
- [22] J. von Neumann and O. Morgenstern. Theory of Games and Economic Behavior. Princeton University Press, 1944.
- [23] C. H. Papadimitriou and T. Roughgarden. Computing Equilibria in Multi-Player Games. *SODA*, 2005.
- [24] C. H. Papadimitriou. Computing Correlated Equilibria in Multiplayer Games. *STOC*, 2005.
- [25] A. Röllin. Translated Poisson Approximation Using Exchangeable Pair Couplings. *ArXiv Report*, 2006.
- [26] B. Roos. Multinomial and Krawtchouk Approximations to the Generalized Multinomial Distribution. *Theory of Probability and Its Applications*, 46(1):103–117, 2001.

APPENDIX

A Missing Proofs

Proof of lemma 3.6: Without loss of generality assume that $0 < \lambda_1 \leq \lambda_2$ and denote $\delta = \lambda_2 - \lambda_1$. For all $i \in \{0, 1, \dots\}$, denote

$$p_i = e^{-\lambda_1} \frac{\lambda_1^i}{i!} \quad \text{and} \quad q_i = e^{-\lambda_2} \frac{\lambda_2^i}{i!}.$$

Finally, define $\mathcal{I}^* = \{i : p_i \geq q_i\}$.

We have

$$\begin{aligned} \sum_{i \in \mathcal{I}^*} |p_i - q_i| &= \sum_{i \in \mathcal{I}^*} (p_i - q_i) \leq \sum_{i \in \mathcal{I}^*} \frac{1}{i!} (e^{-\lambda_1} \lambda_1^i - e^{-\lambda_1 - \delta} \lambda_1^i) \\ &= \sum_{i \in \mathcal{I}^*} \frac{1}{i!} e^{-\lambda_1} \lambda_1^i (1 - e^{-\delta}) \\ &\leq (1 - e^{-\delta}) \sum_{i=0}^n \frac{1}{i!} e^{-\lambda_1} \lambda_1^i = 1 - e^{-\delta}. \end{aligned}$$

On the other hand

$$\begin{aligned} \sum_{i \notin \mathcal{I}^*} |p_i - q_i| &= \sum_{i \notin \mathcal{I}^*} (q_i - p_i) \leq \sum_{i \notin \mathcal{I}^*} \frac{1}{i!} (e^{-\lambda_1} (\lambda_1 + \delta)^i - e^{-\lambda_1} \lambda_1^i) \\ &= \sum_{i \notin \mathcal{I}^*} \frac{1}{i!} e^{-\lambda_1} ((\lambda_1 + \delta)^i - \lambda_1^i) \\ &\leq \sum_{i=0}^n \frac{1}{i!} e^{-\lambda_1} ((\lambda_1 + \delta)^i - \lambda_1^i) \\ &= e^\delta \sum_{i=0}^n \frac{1}{i!} e^{-(\lambda_1 + \delta)} (\lambda_1 + \delta)^i - \sum_{i=0}^n \frac{1}{i!} e^{-\lambda_1} \lambda_1^i \\ &= e^\delta - 1. \end{aligned}$$

Combining the above we get the result. ■

Proof of lemma 3.10: For all $i \in \mathcal{I}$ and any choice of values $p_j \in [0, u]$, $j \in \mathcal{I} \setminus \{i\}$, define the function

$$f(x) = \frac{x^2 + A}{x + B}$$

where $A = \sum_{j \in \mathcal{I} \setminus \{i\}} p_j^2$ and $B = \sum_{j \in \mathcal{I} \setminus \{i\}} p_j$; observe that $A \leq uB$.

If $A = B = 0$ then $f(x) = x$ so f achieves its maximum at $x = u$.

If $A \neq 0 \neq B$, the derivative of f is

$$f'(x) = \frac{-A + x(2B + x)}{(B + x)^2}.$$

Denoting by $h(x)$ the numerator of the above expression, the derivative of h is

$$h'(x) = 2B + 2x > 0, \forall x.$$

Therefore, h is increasing which implies that $h(x) = 0$ has at most one root and hence $f'(x) = 0$ has at most one root since the denominator in the above expression for $f'(x)$ is always positive. Note that $f'(0) < 0$ whereas $f'(u) > 0$.

Therefore, there exists a unique $\rho \in (0, u)$ such that $f'(x) < 0, \forall x \in (0, \rho)$, $f'(\rho) = 0$ and $f'(x) > 0, \forall x \in (\rho, u)$, i.e. f is decreasing in $(0, \rho)$ and increasing in (ρ, u) . This implies that

$$\max_{x \in [0, u]} f(x) = \max\{f(0), f(u)\}.$$

But $f(u) = \frac{u^2 + A}{u + B} \geq \frac{A}{B} = f(0)$ since $A \leq uB$. Therefore, $\max_x f(x) = f(u)$.

From the above it follows that, independent of the values of $\sum_{j \in \mathcal{I} \setminus \{i\}} p_i^2$ and $\sum_{j \in \mathcal{I} \setminus \{i\}} p_i$, $f(x)$ is maximized at $x = u$. Therefore, the expression $\frac{\sum_{i \in \mathcal{I}} p_i^2}{\sum_{i \in \mathcal{I}} p_i}$ is maximized when $p_i = u$ for all $i \in \mathcal{I}$. This implies

$$\frac{\sum_{i \in \mathcal{I}} p_i^2}{\sum_{i \in \mathcal{I}} p_i} \leq \frac{|\mathcal{I}|u^2}{|\mathcal{I}|u} \leq u.$$

■

Proof of lemma 3.11: For (a) we have

$$\begin{aligned} \zeta_j &= \sum_{i \in I_j^*} \mathcal{E}[X_i] - \sum_{i \in I_j^*} \mathcal{E}[Y_i] = \sum_{i=1}^{n_j} p_i^j - \sum_{i=1}^{n_j} q_i^j \\ &= \sum_{i=1}^{n_j} \left(\frac{j}{k} + \delta_i^j \right) - \left(m_j \frac{j+1}{k} + (n_j - m_j) \frac{j}{k} \right) \\ &= \sum_{i=1}^{n_j} \delta_i^j - m_j \frac{1}{k}. \end{aligned}$$

To get (b) we notice that

$$\begin{aligned} \sum_{i=1}^{n_j} \delta_i^j - m_j \frac{1}{k} &= S_j - \lfloor S_j k \rfloor \frac{1}{k} \\ &= S_j - (S_j k - (S_j k - \lfloor S_j k \rfloor)) \frac{1}{k} \\ &= (S_j k - \lfloor S_j k \rfloor) \frac{1}{k}. \end{aligned}$$

For (c), (d) and (e), observe that, for all $j_i \in I_j^*, i \in \{1, \dots, n_j\}$,

$$\begin{aligned} \text{Var}[X_{j_i}] &= p_i^j(1 - p_i^j) = p_i^j - (p_i^j)^2 = \mathcal{E}[X_{j_i}] - \left(\frac{j}{k} + \delta_i^j \right)^2 \\ &= \frac{j}{k} + \delta_i^j - \left(\frac{j}{k} + \delta_i^j \right)^2 \end{aligned}$$

$$\begin{aligned} \text{Var}[Y_{j_i}] &= q_i^j - (q_i^j)^2 = \begin{cases} \mathcal{E}[Y_{j_i}] - \frac{(j+1)^2}{k^2}, & i \in \{1, \dots, m_j\} \\ \mathcal{E}[Y_{j_i}] - \frac{j^2}{k^2}, & i \in \{m_j + 1, \dots, n_j\} \end{cases} \\ &= \begin{cases} \frac{j+1}{k} - \frac{(j+1)^2}{k^2}, & i \in \{1, \dots, m_j\} \\ \frac{j}{k} - \frac{j^2}{k^2}, & i \in \{m_j + 1, \dots, n_j\} \end{cases} \end{aligned}$$

Hence,

$$\begin{aligned}
\sum_{i \in I_j^*} \text{Var}[X_i] &= \sum_{i=1}^{n_j} \text{Var}[X_{j_i}] = \sum_{i=1}^{n_j} \mathcal{E}[X_{j_i}] - n_j \frac{j^2}{k^2} - 2 \frac{j}{k} \sum_{i=1}^{n_j} \delta_i^j - \sum_{i=1}^{n_j} (\delta_i^j)^2 \\
&= n_j \frac{j}{k} + \sum_{i=1}^{n_j} \delta_i^j - n_j \frac{j^2}{k^2} - 2 \frac{j}{k} \sum_{i=1}^{n_j} \delta_i^j - \sum_{i=1}^{n_j} (\delta_i^j)^2 \\
&= n_j \frac{j}{k} \left(1 - \frac{j}{k}\right) + \left(1 - \frac{2j}{k}\right) \sum_{i=1}^{n_j} \delta_i^j - \sum_{i=1}^{n_j} (\delta_i^j)^2.
\end{aligned}$$

$$\begin{aligned}
\sum_{i \in I_j^*} \text{Var}[Y_i] &= \sum_{i=1}^{n_j} \text{Var}[Y_{j_i}] = \sum_{i=1}^{n_j} \mathcal{E}[Y_{j_i}] - n_j \frac{j^2}{k^2} - m_j \frac{2j+1}{k^2} \\
&= n_j \frac{j}{k} + m_j \frac{1}{k} - n_j \frac{j^2}{k^2} - m_j \frac{2j+1}{k^2} \\
&= n_j \frac{j}{k} \left(1 - \frac{j}{k}\right) + m_j \frac{1}{k} \left(1 - \frac{2j+1}{k}\right).
\end{aligned}$$

and

$$\begin{aligned}
\sum_{i \in I_j^*} \text{Var}[X_i] - \sum_{i \in I_j^*} \text{Var}[Y_i] &= \\
&= \left(n_j \frac{j}{k} \left(1 - \frac{j}{k}\right) + \left(1 - \frac{2j}{k}\right) \sum_{i=1}^{n_j} \delta_i^j - \sum_{i=1}^{n_j} (\delta_i^j)^2 \right) - \left(n_j \frac{j}{k} \left(1 - \frac{j}{k}\right) + m_j \frac{1}{k} \left(1 - \frac{2j+1}{k}\right) \right) \\
&= \left(1 - \frac{2j}{k}\right) \left(\sum_{i=1}^{n_j} \delta_i^j - m_j \frac{1}{k} \right) + \left(m_j \frac{1}{k^2} - \sum_{i=1}^{n_j} (\delta_i^j)^2 \right) \\
&= \left(1 - \frac{2j}{k}\right) \zeta_j + \left(m_j \frac{1}{k^2} - \sum_{i=1}^{n_j} (\delta_i^j)^2 \right).
\end{aligned}$$

■

Proof of Lemma 3.13: The coupling lemma implies that for any joint distribution on $\{X_i\}_i \cup \{Y_i\}_i$ the following is satisfied

$$\left\| \sum_{i \in \mathcal{M}_1^*(k)} X_i - \sum_{i \in \mathcal{M}_1^*(k)} Y_i \right\| \leq \Pr \left[\sum_{i \in \mathcal{M}_1^*(k)} X_i \neq \sum_{i \in \mathcal{M}_1^*(k)} Y_i \right].$$

A union bound further implies

$$\Pr \left[\sum_{i \in \mathcal{M}_1^*(k)} X_i \neq \sum_{i \in \mathcal{M}_1^*(k)} Y_i \right] \leq \Pr \left[\bigvee_{i \in \mathcal{M}_1^*(k)} X_i \neq Y_i \right] \leq \sum_{i \in \mathcal{M}_1^*(k)} \Pr[X_i \neq Y_i].$$

Hence for any joint distribution on $\{X_i\}_i \cup \{Y_i\}_i$ the following is satisfied

$$\left\| \sum_{i \in \mathcal{M}_1^*(k)} X_i - \sum_{i \in \mathcal{M}_1^*(k)} Y_i \right\| \leq \sum_{i \in \mathcal{M}_1^*(k)} \Pr[X_i \neq Y_i]. \quad (3)$$

Let us now choose a joint distribution on $\{X_i\}_i \cup \{Y_i\}_i$ in which, for all i , X_i and Y_i are coupled in such a way that

$$\Pr[X_i \neq Y_i] \leq \frac{1}{k}.$$

This is easy to do since by construction $|p_i - q_i| \leq \frac{1}{k}$, for all i . Plugging in Formula (3) the particular joint distribution just described yields

$$\left\| \sum_{i \in \mathcal{M}_1^*(k)} X_i - \sum_{i \in \mathcal{M}_1^*(k)} Y_i \right\| \leq \frac{|\mathcal{M}_1^*(k)|}{k} \leq \frac{1}{k^{1-\beta}}.$$

■

Proof of lemma 3.15: For all $i \in \mathcal{I}$ and any choice of values $p_j \in [u, \frac{1}{2}]$, $j \in \mathcal{I} \setminus \{i\}$, define the function

$$f(x) = \frac{x^3(1-x) + A}{(x(1-x) + B)^2}, x \in \left[u, \frac{1}{2}\right]$$

where $A = \sum_{j \in \mathcal{I} \setminus \{i\}} p_j^3(1-p_j)$ and $B = \sum_{j \in \mathcal{I} \setminus \{i\}} p_j(1-p_j)$.

For the sake of the argument let us extend the range of f to $[0, \frac{1}{2}]$. The derivative of f is

$$f'(x) = \frac{A(-2+4x) + x^2(3B+x-4Bx-x^2)}{(B+x-x^2)^3},$$

where note that the denominator is positive for all $x \in [0, \frac{1}{2}]$. Denoting by $h(x)$ the numerator of the above expression, the derivative of h is

$$h'(x) = 4A + (1-2x)x^2 + 2x(x-x^2) + 6Bx(1-2x) > 0, \forall x \in \left[0, \frac{1}{2}\right].$$

Therefore, h is increasing in $(0, \frac{1}{2})$ which implies that $h(x) = 0$ has at most one root in $(0, \frac{1}{2})$ and hence $f'(x) = 0$ has at most one root in $(0, \frac{1}{2})$ since the denominator in the above expression for $f'(x)$ is always positive. Note that $f'(0) < 0$ whereas $f'(\frac{1}{2}) > 0$. Therefore, there exists a unique $\rho \in (0, \frac{1}{2})$ such that $f'(x) < 0, \forall x \in (0, \rho)$, $f'(\rho) = 0$ and $f'(x) > 0, \forall x \in (\rho, \frac{1}{2})$, i.e. f is decreasing in $(0, \rho)$ and increasing in $(\rho, \frac{1}{2})$. This implies that

$$\max_{x \in [u, 1/2]} f(x) = \max\{f(u), f(1/2)\}.$$

Hence, the expression $\frac{\sum_{i \in \mathcal{I}} p_i^3(1-p_i)}{(\sum_{i \in \mathcal{I}} p_i(1-p_i))^2}$ is bounded by

$$\max_{m \in \{0, \dots, |\mathcal{I}|\}} \frac{\frac{m}{16} + (|\mathcal{I}| - m)u^3(1-u)}{(\frac{m}{4} + (|\mathcal{I}| - m)u(1-u))^2}$$

which we further bound by

$$\max_{x \in [0, |\mathcal{I}|]} g(x)$$

where

$$g(x) := \frac{\frac{x}{16} + (|\mathcal{I}| - x)u^3(1-u)}{(\frac{x}{4} + (|\mathcal{I}| - x)u(1-u))^2}.$$

The derivative of g is

$$g'(x) = \frac{4|\mathcal{I}|u(1+u-6u^2+12u^3-8u^4) - (1-2u-16u^3+48u^4-32u^5)x}{(1-2u)^{-1}(4|\mathcal{I}|(1-u)u + (1-2u)^2x)^3},$$

where note that the denominator is positive for all $u \in [0, \frac{1}{2}]$ and the numerator is of the form $A(u) - B(u)x$ where

$$A(u) := 4|\mathcal{I}|u(1+u-6u^2+12u^3-8u^4) > 0, \forall u \in \left(0, \frac{1}{2}\right)$$

and

$$B(u) := 1-2u-16u^3+48u^4-32u^5 > 0, \forall u \in \left(0, \frac{1}{2}\right).$$

Hence, if we take $\zeta := \frac{A(u)}{B(u)}$, $g'(x) > 0$, $\forall x \in (0, \zeta)$, $g'(\zeta) = 0$ and $g'(x) < 0$, $\forall x \in (\zeta, +\infty)$, i.e. g is increasing in $(0, \zeta)$ and decreasing in $(\zeta, +\infty)$. This implies that g achieves its maximum at $x := \zeta$. The maximum value itself is

$$g(\zeta) = \frac{(1 + 2u + 4u^2 - 8u^3)^2}{16|\mathcal{I}|u(1 - u - 4u^2 + 4u^3)}.$$

This concludes the proof. ■

Proof of claim 3.16: From Lemma 3.7 it follows that

$$||TP(\mu_1, \sigma_1^2) - TP(\mu_2, \sigma_2^2)|| \leq \max \left\{ \frac{|\mu_1 - \mu_2|}{\sigma_1} + \frac{|\sigma_1^2 - \sigma_2^2| + 1}{\sigma_1^2}, \frac{|\mu_1 - \mu_2|}{\sigma_2} + \frac{|\sigma_1^2 - \sigma_2^2| + 1}{\sigma_2^2} \right\}.$$

Denoting $J = \{\lfloor k^\alpha \rfloor, \dots, \lfloor k/2 \rfloor - 1\}$, we have that

$$\begin{aligned} \mu_1 - \mu_2 &= \sum_{i \in \mathcal{M}_1^*(k)} \mathcal{E}[X_i] - \sum_{i \in \mathcal{M}_1^*(k)} \mathcal{E}[Y_i] = \sum_{j \in J} \sum_{i \in I_j^*} (\mathcal{E}[X_i] - \mathcal{E}[Y_i]) = \sum_{j \in J} \zeta_j. \\ \sigma_1^2 &= \sum_{j \in J} \sum_{i \in I_j^*} \text{Var}[X_i] = \sum_{j \in J} \left(n_j \frac{j}{k} \left(1 - \frac{j}{k} \right) + \left(1 - \frac{2j}{k} \right) \sum_{i=1}^{n_j} \delta_i^j - \sum_{i=1}^{n_j} (\delta_i^j)^2 \right) \\ &\geq \sum_{j \in J} \left(n_j \frac{j}{k} \left(1 - \frac{j}{k} \right) - n_j \frac{1}{k^2} \right) \\ &\geq \sum_{j \in J} \frac{n_j}{k^2} (j(k-j) - 1). \end{aligned}$$

Similarly

$$\begin{aligned} \sigma_2^2 &= \sum_{j \in J} \sum_{i \in I_j^*} \text{Var}[Y_i] = \sum_{j \in J} \left(n_j \frac{j}{k} \left(1 - \frac{j}{k} \right) + m_j \frac{1}{k} \left(1 - \frac{2j+1}{k} \right) \right) \\ &\geq \sum_{j \in J} \left(n_j \frac{j}{k} \left(1 - \frac{j}{k} \right) \right) \\ &\geq \sum_{j \in J} \frac{n_j}{k^2} j(k-j) \\ &\geq \sum_{j \in J} \frac{n_j}{k^2} (j(k-j) - 1). \end{aligned}$$

Finally,

$$\sigma_1^2 - \sigma_2^2 = \sum_{j \in J} \left(\sum_{i \in I_j^*} \text{Var}[X_i] - \sum_{i \in I_j^*} \text{Var}[Y_i] \right) = \sum_{j \in J} \left(\left(1 - \frac{2j}{k} \right) \zeta_j + \left(m_j \frac{1}{k^2} - \sum_{i=1}^{n_j} (\delta_i^j)^2 \right) \right),$$

where observe that $\sigma_1^2 - \sigma_2^2 \geq 0$, since

$$\begin{aligned} \left(1 - \frac{2j}{k} \right) \zeta_j + \left(m_j \frac{1}{k^2} - \sum_{i=1}^{n_j} (\delta_i^j)^2 \right) &\geq \left(1 - \frac{2j}{k} \right) \zeta_j + \left(m_j \frac{1}{k^2} - \frac{1}{k} \sum_{i=1}^{n_j} \delta_i^j \right) \\ &= \left(1 - \frac{2j}{k} \right) \zeta_j - \frac{1}{k} \zeta_j = \\ &= \left(1 - \frac{2j+1}{k} \right) \zeta_j \geq 0. \end{aligned}$$

We proceed to bound each of the terms $\frac{|\mu_1 - \mu_2|}{\sigma_1}$, $\frac{|\sigma_1^2 - \sigma_2^2| + 1}{\sigma_1^2}$, $\frac{|\mu_1 - \mu_2|}{\sigma_2}$ and $\frac{|\sigma_1^2 - \sigma_2^2| + 1}{\sigma_2^2}$ separately. We have

$$\begin{aligned}
\frac{|\mu_1 - \mu_2|}{\sigma_1} &= \frac{\mu_1 - \mu_2}{\sigma_1} \leq \frac{\sum_{j \in J} \zeta_j}{\sqrt{\sum_{j \in J} \frac{n_j}{k^2} (j(k-j) - 1)}} \\
&\leq \frac{k \sum_{j \in J} \zeta_j}{\sqrt{\sum_{j \in J: n_j \geq 1} (j(k-j) - 1)}} \\
&\leq \frac{Z}{\sqrt{\sum_{j \in J: n_j \geq 1} (j(k-j) - 1)}} \quad \text{where } Z = |\{j \in J | n_j \geq 1\}| \\
&= \frac{1}{\sqrt{\sum_{j \in J: n_j \geq 1} \left(\frac{j(k-j) - 1}{Z^2} \right)}} \\
&\leq \frac{1}{\sqrt{\sum_{j=\lfloor k^\alpha \rfloor}^{\lfloor k^\alpha \rfloor + Z - 1} \left(\frac{j(k-j) - 1}{Z^2} \right)}} \\
&\leq \frac{1}{\sqrt{\min_{1 \leq Z \leq k/2 - \lfloor k^\alpha \rfloor} \left\{ \sum_{j=\lfloor k^\alpha \rfloor}^{\lfloor k^\alpha \rfloor + Z - 1} \left(\frac{j(k-j) - 1}{Z^2} \right) \right\}}}.
\end{aligned}$$

Note that

$$\sum_{j=j_1}^{j_2+Z-1} (j(k-j) - 1) = \frac{1}{6} Z(-7 - 6j_1^2 + 6j_1(1+k-Z) + 3k(-1+Z) + 3Z - 2Z^2). \quad (4)$$

For $j_1 = \lfloor k^\alpha \rfloor$, let us define

$$F(Z) := \frac{1}{Z^2} \sum_{j=j_1}^{j_2+Z-1} (j(k-j) - 1) \equiv \frac{1}{6Z} (-7 - 6j_1^2 + 6j_1(1+k-Z) + 3k(-1+Z) + 3Z - 2Z^2)$$

The derivative of F is

$$F'(Z) = \frac{7 + 6j_1^2 + 3k - 6j_1(1+k) - 2Z^2}{6Z^2} = \frac{7 + 6\lfloor k^\alpha \rfloor^2 + 3k - 6\lfloor k^\alpha \rfloor(1+k) - 2Z^2}{6Z^2} < 0, \text{ for large enough } k.$$

Hence F is decreasing in $[1, k/2 - \lfloor k^\alpha \rfloor]$ so it achieves its minimum at $Z = k/2 - \lfloor k^\alpha \rfloor$. The minimum itself is

$$F(k/2 - \lfloor k^\alpha \rfloor) = \frac{-14 + 6\lfloor k^\alpha \rfloor - 4\lfloor k^\alpha \rfloor^2 - 3k + 4\lfloor k^\alpha \rfloor k + 2k^2}{6k - 12\lfloor k^\alpha \rfloor} = \Omega(k).$$

Hence,

$$\frac{|\mu_1 - \mu_2|}{\sigma_1} = O(k^{-1/2}).$$

Similarly, we get

$$\frac{|\mu_1 - \mu_2|}{\sigma_2} = O(k^{-1/2}).$$

It remains to bound the terms $\frac{|\sigma_1^2 - \sigma_2^2| + 1}{\sigma_1^2}$ and $\frac{|\sigma_1^2 - \sigma_2^2| + 1}{\sigma_2^2}$. We have

$$\begin{aligned}
\frac{|\sigma_1^2 - \sigma_2^2| + 1}{\sigma_1^2} &\leq \frac{1 + \sum_{j \in J} (\zeta_j + m_j \frac{1}{k^2})}{\sum_{j \in J} \frac{n_j}{k^2} (j(k-j) - 1)} \\
&\leq \frac{1 + k \sum_{j \in J} (k \zeta_j) + M}{\sum_{j \in J} n_j (j(k-j) - 1)} \quad (\text{where } M := \sum_j m_j) \\
&\leq \frac{1 + kZ + M}{\sum_{j \in J} n_j (j(k-j) - 1)} \quad (\text{where } Z = |\{j \in J | n_j \geq 1\}|) \\
&= \frac{kZ}{\sum_{j \in J} n_j (j(k-j) - 1)} + \frac{1 + M}{\sum_{j \in J} n_j (j(k-j) - 1)}
\end{aligned}$$

The second term of the above expression is bounded as follows

$$\begin{aligned}
\frac{1 + M}{\sum_{j \in J} n_j (j(k-j) - 1)} &\leq \frac{1}{\sum_{j \in J} n_j (j(k-j) - 1)} + \frac{M}{\sum_{j \in J} n_j (j(k-j) - 1)} \\
&\leq \frac{1}{\sum_{j \in J} n_j (j(k-j) - 1)} + \frac{M}{\sum_{j \in J} n_j (j(k-j) - 1)} \\
&\leq \frac{1}{\lfloor k^\alpha \rfloor (k - \lfloor k^\alpha \rfloor) - 1} + \frac{1}{\sum_{j \in J} \frac{m_j}{M} (\lfloor k^\alpha \rfloor (k - \lfloor k^\alpha \rfloor) - 1)} \\
&\leq \frac{2}{\lfloor k^\alpha \rfloor (k - \lfloor k^\alpha \rfloor) - 1} = O(k^{-1-\alpha}).
\end{aligned}$$

The first term is bounded as follows

$$\begin{aligned}
\frac{kZ}{\sum_{j \in J} n_j (j(k-j) - 1)} &\leq \frac{kZ}{\sum_{j \in J: n_j \geq 1} (j(k-j) - 1)} \\
&\leq \frac{1}{\sum_{j=\lfloor k^\alpha \rfloor}^{\lfloor k^\alpha \rfloor + Z - 1} \left(\frac{j(k-j) - 1}{kZ} \right)} \\
&\leq \frac{1}{\min_{1 \leq Z \leq k/2 - \lfloor k^\alpha \rfloor} \left\{ \sum_{j=\lfloor k^\alpha \rfloor}^{\lfloor k^\alpha \rfloor + Z - 1} \left(\frac{j(k-j) - 1}{kZ} \right) \right\}}.
\end{aligned}$$

From above

$$\sum_{j=j_1}^{j_2+Z-1} (j(k-j) - 1) = \frac{1}{6} Z (-7 - 6j_1^2 + 6j_1(1+k-Z) + 3k(-1+Z) + 3Z - 2Z^2). \quad (5)$$

For $j_1 = \lfloor k^\alpha \rfloor$, let us define

$$G(Z) := \frac{1}{Zk} \sum_{j=j_1}^{j_2+Z-1} (j(k-j) - 1) \equiv \frac{1}{6k} (-7 - 6j_1^2 + 6j_1(1+k-Z) + 3k(-1+Z) + 3Z - 2Z^2)$$

The derivative of G is

$$G'(Z) = \frac{3 - 6j_1 + 3k - 4Z}{6k} = \frac{3 - 6\lfloor k^\alpha \rfloor + 3k - 4Z}{6k} > 0, \text{ for large enough } k \text{ since } Z \leq k/2.$$

Hence G is increasing in $[1, k/2 - \lfloor k^\alpha \rfloor]$ so it achieves its minimum at $Z = 1$. The minimum itself is

$$G(1) = \frac{-6 - 6\lfloor k^\alpha \rfloor^2 + 6\lfloor k^\alpha \rfloor k}{6k} = \Omega(k^\alpha).$$

Hence,

$$\frac{kZ}{\sum_{j \in J} n_j (j(k-j) - 1)} = O(k^{-\alpha}),$$

which together with the above implies

$$\frac{|\sigma_1^2 - \sigma_2^2| + 1}{\sigma_1^2} = O(k^{-\alpha})$$

Similarly, we get

$$\frac{|\sigma_1^2 - \sigma_2^2| + 1}{\sigma_2^2} = O(k^{-\alpha}).$$

Putting everything together we get that

$$\left| TP(\mu_1, \sigma_1^2) - TP(\mu_2, \sigma_2^2) \right| \leq O(k^{-\alpha}) + O(k^{-1/2}).$$

■